Some challenges with IP multicast deployment

Submission to IAB Design Expectations vs Deployment Reality in Protocol Development Workshop 2019

Lucas Pardue - lucaspardue.24.7@gmail.com

This submission was written intentionally brief and sarcastic. It may have failed on both counts.

Design expectations: all the benefits of IP broadcast with none of the draw backs.
Assumptions: uncoordinated deployments would figure out routing and be friendly.
Deployment reality: "no one deploys IP multicast"
Evidence: The Wikipedia article on the Mbone (multicast backbone) uses the past tense.

Realities of administrative deployment

A common misconception about IP multicast is that it is not deployed for some definition of "real traffic". That although it has some legitimate usage in localised "private networks", WAN and internet-scale deployments are impractical.

It unfortunately seems that many opinions on IP multicast stem from experience with early experimentation of IPv4 any-source multicast (ASM). In that period, the routing of packets between domains was not well defined. The Internet-Draft draft-acg-mboned-deprecate-interdomain-asm provides a good historical overview, it highlights the popularity of RFC 7761 PIM-SM for routing and the challenges this has between domains. It also links to RFC 3618 - Multicast Source Discovery Protocol (MSDP) and its cousin RFC 4611 MSDP Deployment Scenarios, showing deployment scenarios cover both one-to-many distribution, and ad-hoc come-and-go distribution.

To quote the I-D: "To this day, there is no IETF Proposed Standard level interdomain solution for IPv4 ASM multicast because MSDP was the "best" component for the interdomain discovery problem, and it stayed Experimental.". This "best" solution has a poor security story with little way to authenticate senders and protect networks from BUM traffic.

In practice: multicast typically works well within domains, or depends on heavy monitoring and control of ingress traffic. The dream of uncoordinated deployment fell on its sword. We might call these domain multicast islands. In a sea of unicast capability, endpoints are stranded without boats. This can be overcome with coordinated handoffs: formal/static relationships such as tunneling using GRE, or adhoc relationships using some relay infrastructure such as in RFC 7450 Automatic Multicast Tunnelling (AMT). Unfortunately, AMT has its own story of design vs. reality. A brief survey of relays was presented at IETF 101.

Realities of hardware and software

There could be arguments that the IP multicast design made assumptions about the behaviour of the network at Layer 1 and 2. I don't know because I wasn't alive. What I can be more sure about is that the characteristics of diverse Layer 1 protocols and the hardware they are deployed across have affected IP multicast.

In practical terms, network equipment has constraints and attempting to route large multicast deployments take up valuable space and performance. Some of the problems with IPv4 ASM are solved with source-specific multicast (SSM) and IPv6. The deployment reality of IPv6 SSM has a shared fate with IPv6 unicast. Some of the improvements fit better within the constraints of some network equipment. However, slow adoption of technologies has hampered the deployment reality. A perpetual cycle of low demand for services, causing low uptake of required technology, causing low provision of services ensues. For example, the slow rollout of SSM subscription protocols in the form of IGMPv3 and MLDv2.

One area that has overcome this cycle is the vertical market, predominantly IPTV style one-to-many deployments from "triple play" telecomms vendors that provide television, internet and telephone service. These vendors are empowered by their ownership of the domain, along with sending, routing and receiving hardware. The actual propensity of IP multicast by such vendors varies by market. One interesting unintended consequence of TV being major a multicast deployment driver is that consumer premises equipment (e.g. a Set Top Box or home router), which is sensitive to cost, optimises for the use case. The network fabric can often only support subscription to a handful of concurrent multicast groups, each one representing a TV channel.

There are further particular challenges on the consumer side, even in a triple play deployment. If we postulate that the success a deployment is "end-to-end", then one end is the user. Therefore, internet technologies and network services must provide affordance for user expectations when engaging with them. Two aspects of this are discoverability and seamlessness. To many, the internet *is* the web or, worse still, some portal site or social media application. Multicast services in their present form do not fit the Web PKI system, giving users (and their agents) little ability to discover them or switch. Some efforts exist to convert multicast content into a web-friendly delivery mode such as HTTP, for example, the work of the Digital Video Broadcasting Project Multicast ABR group. However, these efforts have trouble satisfying the authenticity and trust requirements that user agents continue to raise the bar on. Some new techniques such as HTTP Signed Exchanges may mitigate this but we should also be cognisant that the internet isn't HTTP alone and that it is not always feasible or practical to deliver content that way. A good example here is the design and development of WebRTC, and how this has been incorporated in the Web trust model.

Returning to hardware challenges, multicast over WiFi is troubled. The I-D draft-ietf-mboned-ieee802-mcast-problems provides an excellent overview of "well-known issues with multicast have prevented the deployment of multicast in 802.11 and other local-area wireless environments."

There must always be an obligatory mention of 5G. Or to be more useful, the deployment of IP multicast on cellular service. Challenges here involve profiling to provide multicast-like services, such as 3GPP Multimedia Broadcast Multicast Service. Constraining IP capabilities in this way increases friction between the delivery of services end-to-end. The quality and capability of hardware and software is likely to suffer the same problems already shown.

Why does anyone care?

The classic poster child is video streaming will break the internet by causing periods of legitimate high traffic volume that overwhelms network segments. Content Distribution Networks (CDN) may go someway to solving this problem but the constraints of individual Network Interface Cards mean that many thousands of machines would be required in order to deliver the aggregate peak bandwidth. Deploying and managing such a CDN will have administrative and security challenges, and lets not forget the impact on the climate. In contrast, IP packet replication can achieve the same result with less problems - if only it worked.

So lets turn the question on its head. If the current challenges with IP multicast deployment on the heterogenous Internet could be solved, what compelling case is there to grow CDN networks for peak traffic flows?

In multicast is a dirty word, perhaps BIER (Bit-indexed Explicit Replication) is the dissociation that we need.